REPORT DOCUMENTATION PAGE

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
19-01-2012	final	

4. TITLE AND SUBTITLE
DISTRIBUTED DECISIONS FOR NETWORKED SYSTEMS

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA9550-08-0085

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Munther A. Dahleh

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
LIDS/EECS-MIT

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AFOSR
875 N. Randolph St.
Suite 325
Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-OSR-VA-TR-2012-0218

12. DISTRIBUTION/AVAILABILITY STATEMENT
A

13. SUPPLEMENTARY NOTES

14. ABSTRACT
We summarize three problems. In the first problem, we study the impact of communication constraints on delay-critical problems of distributed computation. We consider networks of agents, each having access to some partial information, which aim at computing some application-specific function of the global information. Computation has to be completely distributed, i.e., agents can rely on the local observations only, while iteratively processing the available information, and communicating through digital noisy channels. As large delays can be detrimental for the overall system performance, it is critical to design distributed algorithms which perform such computation in the quickest possible way. We present tight lower bounds on the computation delay that capture the connectivity of the network as well as the channel characteristics.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
					19b. TELEPHONE NUMBER (Include area code)

Reset

DISTRIBUTED DECISIONS FOR NETWORKED SYSTEMS

FA9550-08-1-0085

Munther A. Dahleh

LIDS/EECS-MIT

Abstract

We summarize three problems. In the first problem, we study the impact of communication constraints on delay-critical problems of distributed computation. We consider networks of agents, each having access to some partial information, which aim at computing some application-specific function of the global information. Computation has to be completely distributed, i.e., agents can rely on the local observations only, while iteratively processing the available information, and communicating through digital noisy channels. As large delays can be detrimental for the overall system performance, it is critical to design distributed algorithms which perform such computation in the quickest possible way. We present tight lower bounds on the computation delay that capture the connectivity of the network as well as the channel characteristics.

In the second problem, we consider the performance of a class of combinatorial optimization problems with dynamic constraints. Such problems arise in applications involving planning of dynamic systems such as unmanned vehicles. An example of such a class is the traveling salesman problem in which the salesman is a vehicle tasked to pass through a finite number of points in the output space in minimal time. We characterize the performance of such problems when the points are sampled uniformly in the output space of the dynamic system and when the number of such points grows. We provide tight lower bounds on this performance which depends on the number of points and the local controllability properties of the dynamic system and then characterize algorithms that perform within a constant of the lower bound.

In the third problem, we consider the stochastic shortest path problem with limited-side information whereby such information is part of the decision variables. In this context, suppose the agent has limited resources with which to gather information about the edge delays in advance of its travel -- for example, the agent may have a budget for the purchase of sensors, and

it can spend its budget on a handful of high-quality sensors, many low-quality sensors, or some mixture of the two types. In the context of shortest-path optimization, we seek to address the question: what types of sensors should the agent purchase, and on which edges should the agent place the sensors to best improve its overall performance? More generally, we seek to provide a simple, intuitive framework for studying decision-making under limited information, and we seek to provide algorithms that (sub)optimally allocate information resources (such as sensors or bandwidth) to best improve the agent's performance. We also plan to provide intuitive bounds that relate the agent's performance to the amount of information it receives. Finally, as part of the study, we will focus an application of the framework to stochastic shortest-path optimization in order to gain insight into the structures and challenges of such problems.

Status/Progress

Distributed Computation: We consider networks of agents, each having access to some partial information, which aim at computing some application-specific function of the global information. Computation has to be completely distributed, i.e., agents can rely on the local observations only, while iteratively processing the available information, and communicating through digital noisy channels. In our previous work, we utilized Information Theoretic inequalities to derive an algorithm-independent lower bound on the computation time. The bound is a function of the uncertainty in the function to be estimated, via its differential entropy, and the desired accuracy level, as specified by the mean square error criterion. We showed that the lower bound is inversely proportional to the "conductance' of the network, where conductance quantifies the information flow ``bottle-neck" in the network and hence captures the effect of the topology and capacities. This work is reported in [1,2].

In our recent work reported in [3], we provide information-theoretic bounds on the exponential decay rate of the error made in the distributed computation of real-valued functions over noisy networks. Our bounds improve on those proven in our earlier work [2], and show that the worst cut-set capacity is not in general a sufficient measure of performance for delay-critical systems. Indeed, when the communication channels are not deterministic, the decay rate of any moment of the error is bounded away from the min-cut capacity of the network. This is due to the effect of atypical channel realization, which, despite their asymptotically vanishing probability, have a strong impact on the performance. The problem is connected to that of estimating the error exponent of communication channels with feedback.

While the aforementioned bounds are algorithm-independent, in our current research, we are exploring the case in which it is not possible to fully design the dynamics of the algorithm, or when the agents have memory or computation complexity constraints.

Combinatorial Optimization over Dynamics: In this work, we consider the performance of a class of combinatorial optimization problems with dynamic constraints. Such problems arise in applications involving planning of dynamic systems such as unmanned vehicles. As an example, the *traveling salesman problem* (TSP) for a controllable unmanned vehicle can be stated as follows: find a control strategy for a controllable vehicle in order to traverse a closed path through a fixed number of output locations (tasks) in minimum time. This problem is typically solved in two stages: the first is a static traveling salesman problem which finds a closed path through the tasks with minimum length, and the second is finding a control strategy to navigate through the solution of the first stage. The resulting solutions may be far from optimal as such decompositions ignore the system's natural dynamical behavior. Since the exact solutions of this problem can be very hard, we seek to compute the cost when the locations of the tasks are randomly selected in the output space and number of such tasks grows. These results provide generic bounds that capture important properties of both the combinatorial problem as well as the underlying dynamic system. They also provide algorithms that are simple to implement and are order optimal.

We present algorithms for the TSP for controlled dynamic system that, if the points to be visited are randomly sampled from a uniform distribution, is order optimal. In other words, it produces an output trajectory the expected duration of which scales within a constant factor of the optimum, asymptotically in the number of points. We show that the optimal bounds depend on the number of points raised to an exponent directly related to output controllability indices similar to the Kroneckar indices arising in linear systems.

The approach taken in this paper is quite general as we present a class of functionals for which the average scaling behavior can be computed. We refer to these functionals as Quasi-Euclidean functionals and we demonstrate that they include the classical problems of traveling salesman, bipartite matching, and minimum spanning trees over dynamic constraints. This work is reported in [6].

A related problem to the traveling salesman problem is the *Dynamic Traveling Repairman Problem (DTRP)*, where the points to be visited are not known a priori but are generated over time by an exogenous stochastic process. In this context, it is desired to minimize the expected time between the generation of a point and the time it is visited (and possibly repaired) by the output of the dynamic system. We present algorithms that guarantee that the expected waiting time for a customer scales within a constant factor of the optimum as a function of the point generation rate. This work has been reported in [4,5,6].

Stochastic Shortest Path with Side Information: In the stochastic shortest-path optimization setting, the agent's sensors represent a communication channel between the edge delays and the agent, and its choice of path represents its decision based on its information. Traditionally, the quality of a communication channel is measured by the accuracy with which an agent can

estimate an event from the other end of the channel. However, this traditional measure for quality does not account for how the agent actually applies the information it receives, and so while the agent's information resources may be allocated to minimizing estimation error, they may not best allocated with respect to the agent's performance.

In the general framework of decision-making under limited information, a decision is caste as an optimization over all possible decisions of the agent, and it is assumed that the objective to be optimized (the performance) is a function of both the decision and some random event, thus making the objective random as well. In the case where the agent has no information about the random event other than its probability distribution, the agent simply chooses the decision with the best average performance (the average-optimal solution). On the other hand, if the agent has some information about the random event, it can use this information to estimate the objective and determine the average-optimal decision subject to this information As one would expect, the agent's performance improves when it has any amount of information, but the agent can improve its performance further by optimizing its information resources in advance.

In traditional information theory, an allocation of information resources corresponds to an optimization over all channels of a certain class to yield a minimal bit-error rate subject to a fixed random source (rather than the converse of optimizing the source for a fixed channel). We modify this abstract notion of resource allocation by considering an optimization of the agent's performance instead. In both settings, the optimization over a class of channels really corresponds to an optimization over joint distributions between the input and output of the channel subject to some constraint on the joint distributions (an information constraint); the "bigger" this constraint set, the better the resulting performance of the agent. We parameterize a class of constraint sets by a scalar that we term the capacity, and we show how the agent's performance varies with this capacity.

The proposed framework encapsulates the study of **optimal control under limited information**. In this special case, our framework would correspond to the derivation of control policy that is the mapping of information to a decision. If the information includes the system's state, the control policy would be a closed-loop policy. Ultimately, the control policy is a function of the information available to it, and therefore the optimization of the information resources must be performed alongside the derivation of the control law itself. The latter is the focus of our work.

In the context of shortest-path optimization, the decision space for the agent is the set of paths in the graph, and its performance is taken over the lengths of the paths it takes. Finally, we define the capacity of the agent's information as the variance of its estimate of the edge delays.

The basic framework for optimizing agent performance as well as the information structure is established, and some results concerning the application of the framework to linear programs (which includes shortest-path optimization) have been developed. For stochastic shortest-path optimization, we designed a parameterized set of information constraints that yields a useful notion of capacity, and we computed bounds on the agent's performance improvement subject to a given capacity. We also developed an algorithm for (sub)optimally allocating resources without resorting to the use of techniques that only provide lower bounds for the performance, implying that the information-resource allocation provided by the algorithm is truly (sub)optimal [1].

In order to gain an intuitive understanding of information-resource allocation in shortest-path optimization, we compared several allocation schemes on several graph topologies. In some specific cases, particularly when the weights are iid random variables, it was found that if the agent concentrated its information resources along a single path, it achieved better performance than if it allocated resources across many disjoint paths. It was also determined that the agent's allocation should be spread along the single path rather than concentrated along a few edges of the path [1] (i.e., it should use many low-quality sensors rather than a handful of high-quality sensors). Such results offer some insight into the non-obvious aspects information-resource allocation and may further serve as a basis for future heuristics.

We also successfully applied the framework to two other optimization problems of interest. The first is the case of multiple agents using a road system simultaneously, where we computed the benefit of information to the social welfare of the group, and the second is to the production of energy, where we computed the benefit to the energy supplier of having a consumer forecast [2]

Currently, we are developing tighter lower bounds for the agent's average performance, and we are examining an extension of the framework to a variation sequential learning where the actions of some are observed by others. In this latter study, we seek to define and quantify a notion of information and determine how much information is required for sequence to result in a perfect asymptotic decision-making.

While traditional techniques for optimizing sensing and communication resources are appropriate for situations where either estimation accuracy is paramount or the application of the information is unknown, such techniques may not be appropriate in strategic planning environments where it may be well known how the information will be used. For example, our framework may be useful in optimally determining where to send UAVs to gather reconnaissance information or where to place sensors within a traffic system in order to best improve traffic flow. Furthermore,

by providing analytic performance bounds for a given amount of capacity, one can roughly balance the cost of gathering information to the benefit of applying information.

At the technical level, the work provides new, simple techniques for computing (sub)optimal allocations of resources and analytic bounds relating capacity and performance, which includes new methods for computing bounds for order statistics.

Finally, the application of the framework to shortest-path optimization is significant in its own right, not only as an application to traffic systems or network routing, but because shortest-path optimization is the basis for other optimization techniques (such as dynamic programming), in turn providing a ready basis for the study of other related problems.

Personnel Supported During Duration of Grant
Giacomo Como   Post doc
Ola Ayaso   Graduate Students
Slieman Itani   Graduate Student
Michael Rinehart   Graduate Student/Post doc
Munther A. Dahleh   Professor, MIT

Publications

1. O. Ayaso, D. Shah and M.A. Dahleh. *Counting Bits for Distributed Function Computation*, ISIT, July 2008.

2. O. Ayaso, D. Shah and M.A. Dahleh. *Information Theoretic Bounds for Distributed Computations*, accepted for publication in IEEE Trans., Information Theory.

3. G. Como and M. Dahleh. *Lower bounds on the estimation error in problems of distributed computation*," in Proc. of Information Theory and Applications Workshop, (San Diego, CA), 2009.

4. S. Itani, E. Frazzoli and M. A. Dahleh. *Travelling Salesperson Problems for Dynamic Systems*, IFAC, Seoul, 2008.

5. S. Itani, E. Frazzoli and M. A. Dahleh. *Traveling Salesperson-Type Problems for Dynamic Systems*, submitted to IEEE Trans. on Automatic Control.

6. S. Itani, M. A. Dahleh and E. Frazzoli. *Subadditive Functionals and Dynamic Systems*, submitted to IEEE Trans. on Automatic Control.

7. Michael Rinehart, Munther Dahleh. *The Value of Side Information in Shortest Path Optimization*. Accepted for publication in IEEE Tranactions on Automatic Control, 2010.

8. Michael Rinehart, Munther Dahleh. *The Value of Side Information in Network Flow Optimization*. Submitted to Control and Systems Letters, 2010.

9. Michael Rinehart, Munther Dahleh. *The Value of Sequential Information in Shortest Path Optimization*. Submitted to Automatica, 2010.

11. Michael Rinehart, Munther Dahleh. *The Value of Sequential Information in Shortest Path Optimization*. Proceedings of the 2010 American Control Conference, 2010.

12. Michael Rinehart, Munther Dahleh. *A Graph Reduction for Bounding the Value of Side Information in Shortest Path Optimization*. Proceedings of the 2010 American Control Conference, 2010.

13. Michael Rinehart, Munther Dahleh. *Shortest Path Optimization under Limited Information*. Proceedings of the 2009 IEEE Conference on Decision and Control, 2009.